arXiv:1603.03252v1 [cs.LO] 10 Mar 2016

Extension of PRISM by Synthesis of Optimal Timeouts in Fixed-Delay CTMC

L'uboš Korenčiak, Vojtěch Řehák, and Adrian Farmadin

Faculty of Informatics, Masaryk University, Brno, Czech Republic
{korenciak, rehak, xfarmad}@fi.muni.cz

Abstract. We present a practically appealing extension of the probabilistic model checker PRISM rendering it to handle fixed-delay continuous-time Markov chains (fdCTMCs) with rewards, the equivalent formalism to the deterministic and stochastic Petri nets (DSPNs). fdCTMCs allow transitions with fixed-delays (or timeouts) on top of the traditional transitions with exponential rates. Our extension supports an evaluation of expected reward until reaching a given set of target states. The main contribution is that, considering the fixed-delays as parameters, we implemented a synthesis algorithm that computes the epsilon-optimal values of the fixed-delays minimizing the expected reward. We provide a performance evaluation of the synthesis on practical examples.

1 Introduction

PRISM [9] is an efficient tool for probabilistic model-checking of stochastic systems such as Markov decision processes (MDPs), discrete-time Markov chains (DTMCs), or continuous-time Markov chains (CTMCs). The PRISM community frequently raises requests to incorporate the possibility to express delays with deterministic durations in a CTMC.[1] The standard PRISM recommendation is to approximate the deterministic durations using a phase-type technique [11] and thus obtaining a CTMC. This works for some models, however there are models for which such approximation can cause either a large error or a state space explosion (see, e.g. [7,2]). However, there is a formalism called fixed-delay CTMCs (fdCTMCs) [4,7,1] that is the requested extension of CTMCs by fixed-delay (fd) events, modeling the deterministic transitions or timeouts. Recent result [1] came up with new synthesis algorithms working directly on fdCTMCs (rather than approximating them with CTMCs). Here we provide the first attempt to experimental evaluation of such synthesis algorithms and show that they are practically applicable. In the following running example we demonstrate the fdCTMC semantics as well as the parameters and objectives of the synthesis.

Example 1. The figure bellow depicts fdCTMC of a slightly modified model of dynamic power management of a Fujitsu disk drive taken from the PRISM case studies[2] [13]. The disk has three modes *idle*, *busy*, and *sleep*. In the *idle* and *sleep* modes the disk receives

[1] http://www.prismmodelchecker.org/manual/FrequentlyAskedQuestions/PRISMModelling#det_delay

[2] http://www.prismmodelchecker.org/casestudies/power_ctmc3.php

requests, in the *busy* mode it also serves them. The disk is equipped with a bounded buffer, where it stores requests when they arrive. The requests arrive with an exponential inter-arrival time of rate 1.39 and increase the current size of the buffer. The requests are served in an exponential time of rate 12.5, what decreases the buffer size. Note that restricting the model to the *idle* and *busy* modes only, we obtain a CTMC model of an M/M/1/n queue.

Moreover, the disk can move from the *idle* mode to the *sleep* mode where it saves energy. Switching of the disk to the *sleep* mode is driven by timeout. This is modeled by an fd event f_1 moving the state from $(idle, 0)$ to $(sleep, 0)$ when the disk is steadily idle for a specified amount of time (e.g. 1 second). The disk is woken up by another timeout modeled by an fd event f_2, which is active in all *sleep* states. After staying in the *sleep* mode for, e.g. 2 seconds, f_2 changes the state according to the dashed arrows.

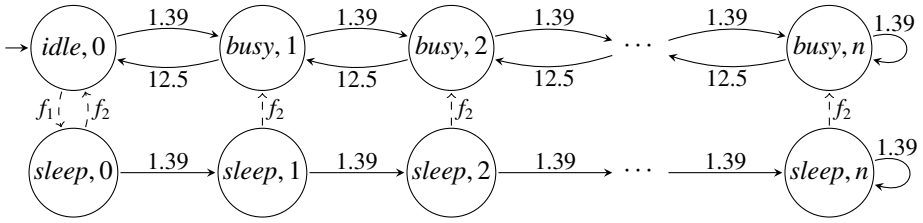

Additionally, every state is given a rate cost that specifies an amount of energy consumed per each second spent there. Optionally, an impulse cost can be specified, e.g., say that the change from $(idle, 0)$ to $(sleep, 0)$ consumes 0.006 energy units instantaneously. Now, one might be interested in how much energy on average is consumed before emptying the buffer, i.e. to compute the expected energy consumed until reaching target that is a new successor of $(busy, 1)$ instead of the initial state $(idle, 0)$. But, being a developer of the disk, can we set better timeouts for f_1 and f_2? Hence, we consider timeouts as parameters and synthesize them in order to minimize the expected amount of consumed energy.

Our Contribution is as follows. 1. We provide an extension of the PRISM language and of the internal data structures to support specification of fdCTMC with impulse and rate costs (or equivalently rewards). Hence, our version of PRISM is now ready for other experiments with fdCTMC algorithms including the possibility to support model-checking options as for CTMCs and DTMCs. 2. We added an evaluation of expected reward until reaching a given set of target states. 3. We analyzed the synthesis algorithm from [1], derived exact formulas and implemented the algorithm. 4. Additionally, we accelerated the implementation by few structural changes, that significantly improved the running time and the space requirements of the synthesis implementation. 5. We provide a performance evaluation proving that current implementation is practically applicable to a complex model from the PRISM case-study.

Related Work There are many papers that contain models with fd events suitable for synthesis such as deterministic durations in train control systems [15], time of server rejuvenation [3], timeouts in power management systems [13], etc. Some of the models already contain specified impulse or rate costs.

In [14] authors compute the optimal value of webserver timeout using impulse and rate costs. The implementation can dynamically change the optimal value of timeout based on the current inter-arrival times of requests. It works on the exact fdCTMC model and cannot be easily applied to the more general fdCTMC models our implementation can handle.

The formalism of deterministic and stochastic Petri nets (DSPNs) is equivalent to fdCTMCs. DSPNs have been extensively studied and many useful results are directly applicable to fdCTMCs. To the best of our knowledge the synthesis of fd events has not been studied for DSPNs. The most useful tools for DSPNs are ORIS [6] and TimeNET [16].

There was also an option to implement the synthesis algorithm as an extension of ORIS. However, PRISM is much more used in practice and contains solution methods for MDPs, that we needed for our implementation. Thus, we decided to implement the synthesis into PRISM, even thought we had to extend the PRISM language and data structures. Therefore, the ORIS and TimeNET algorithms can be now reimplemented for fdCTMCs in PRISM easily, exploiting its efficient symbolic structures and algorithms for CTMCs or MDPs.

In the rest of the paper we first formally define the fdCTMC and explain the extension of PRISM language. Then we discuss the implemented algorithms and the performance results.

2 Preliminaries

We use $\mathbb{N}_0$, $\mathbb{R}_{\geq 0}$, and $\mathbb{R}_{>0}$ to denote the set of all non-negative integers, non-negative real numbers, and positive real numbers, respectively. Furthermore, for a countable set A, we denote by $\mathcal{D}(A)$ the set of discrete probability distributions over A, i.e. functions $\mu : A \to \mathbb{R}_{\geq 0}$ such that $\sum_{a \in A} \mu(a) = 1$.

Definition 1. *A* fixed-delay CTMC *(fdCTMC) C is a tuple $(S, Q, F, A, N, d, s_{in})$ where*

- *S is a finite set of states,*
- *$Q : S \times S \to \mathbb{R}_{\geq 0}$ is a rate matrix,*
- *F is a finite set of fixed-delay (fd) events,*
- *$A : S \to 2^F$ assigns to each state s a set of active fd events in s,*
- *$N : S \times F \to \mathcal{D}(S)$ is the successor function, i.e. assigns a probability distribution specifying the successor state to each state and fd event that is active there,*
- *$d : F \to \mathbb{R}_{>0}$ is a delay vector that assigns a positive delay to each fd event,*
- *$s_{in} \in S$ is an initial state.*

Note that fdCTMC C with empty set of fd events is a CTMC. The fdCTMC formalism can be understood as a stochastic event-driven system, i.e. the amount of time spent in each state and the probability of moving to the next state is driven by the occurrence of events. In addition to the fd events of F, there is an *exponential event* $\mathcal{E}$ that is active in all states s where $\sum_{s' \in S} Q(s, s') > 0$. During an execution of an fdCTMC all active events keep one *timer*, that holds the remaining time until the event occurs. The execution starts in the state s_{in}. The timer of each fd event f in $A(s_{in})$ is set to $d(f)$. The timer of the exponential event is set randomly according to the exponential distribution with a rate

$\sum_{s' \in S} Q(s_{in}, s')$. The event e with least[3] timer value t occurs and causes change of state. In case e is an fd event, the next state is chosen randomly according to the distribution $N(s_{in}, e)$, otherwise e is an exponential event and the probability of choosing s as a next state is $Q(s_{in}, s) / \sum_{s' \in S} Q(s_{in}, s')$. In the next state s, the timers of all newly active fd events (i.e. $A(s) \setminus A(s_{in})$), the occurred event e, and the exponential event are set in the same way as above. Observe that the timers of the remaining active fd events decreased by time t spent in the previous state. The execution then proceeds in the same manner.

We illustrate the definition on the fdCTMC model from Example 1. The execution starts in $(idle, 0)$. The events f_1 and $\mathcal{E}$ are active and their timers are set to 1 and e.g. 1.18, respectively. Hence, after 1 second f_1 occurs and changes the state to $(sleep, 0)$ with probability 1. The timers of newly active event f_2 and $\mathcal{E}$ are set to 2 and e.g. 1.5, respectively. Now, $\mathcal{E}$ occurs and changes the state to $(sleep, 1)$. Here f_2 is still active and thus its timer holds the original value subtracted by the time spent in $(sleep, 0)$, i.e. $2 - 1.5 = 0.5$. The timer of the exponential event is set, etc.

A *run* of the fdCTMC is an infinite sequence $(s_0, e_0, t_0)(s_1, e_1, t_1) \dots$ where $s_0 = s_{in}$ and for each $i \in \mathbb{N}_0$ it holds that $s_i \in S$ is the i-th visited state, $e_i \in \{\mathcal{E}\} \cup F$ is the event that occurred in s_i, and $t_i \in \mathbb{R}_{\geq 0}$ is the time spent in s_i. For the formal definition of the semantics of fdCTMC and the probability space on runs see [8].

Total Reward Before Reaching a Target To allow formalization of performance properties we enrich the model in a standard way with rewards or costs (see, e.g. [12]). For an fdCTMC C with a state space S we additionally define a set of target states T, reward rates $\mathcal{R}$, and impulse rewards $\mathcal{I}$. Formally, the target state T is a subset of $S \setminus s_{in}$, $\mathcal{R} : S \to \mathbb{R}_{\geq 0}$ assigns a reward rate to every state, and $\mathcal{I} : S \times (\{\mathcal{E}\} \cup F) \times S \to \mathbb{R}_{\geq 0}$ assigns an impulse reward to every change of state. Now the reward assigned to a run $(s_0, e_0, t_0)(s_1, e_1, t_1) \dots$ is the reward accumulated before reaching a state of T, i.e. $\sum_{i=0}^{n-1} (t_i \cdot \mathcal{R}(s_i) + \mathcal{I}(s_i, e_i, s_{i+1}))$ where $n > 0$ is the minimal index such that $s_n \in T$. We set the reward to infinity whenever there is no such n. The reward of a run can be viewed as a random variable, say $Cost_{C,T,\mathcal{R},\mathcal{I}}$. By $E_{C,T,\mathcal{R},\mathcal{I}}$ (or simply E_C) we denote the expected value of $Cost_{C,T,\mathcal{R},\mathcal{I}}$.

Synthesis Given a delay vector d', let (parametric) fdCTMC $C(d')$ be the fdCTMC C where the delay vector is changed to d'. Our aim is to find a delay vector d such that the expected reward $E_{C(d)}$ is minimal. Formally, given an error bound $\varepsilon > 0$ the synthesis algorithm computes delay vector d, such that $E_{C(d)} \leq Val[C] + \varepsilon$, where $Val[C]$ denotes the optimal reward $\inf_{d'} E_{C(d')}$.

3 PRISM Language and User Interface Extension

Each fdCTMC model file must begin with the keyword `fdctmc`. For the purpose of our synthesis and expected reward implementation, the set of target states has to be specified by label `"target"`, e.g.

```
label "target" = s=2;
```

[3] For the sake of simplicity, when multiple events $X = \{e_1, \dots, e_n\}$ occur simultaneously, the successor is determined by the minimal element of X according to some fixed total order on F.

The exponential event (the matrix Q) is specified the same way as in CTMC models of PRISM. The fd events are local to a module and must be declared immediately after the module name. E.g. the `fdelay f = 1.0` defines the fd event f with delay of a double value 1.0. For an fd event f we specify its set of active states (i.e. $A^{-1}(f)$) and transition kernel (i.e. $N(\cdot, f)$) by PRISM commands where the identifier f is in the arrow. E.g.

```
[L] s=1 --f-> 0.3:(s'=0) + 0.7:(s'=2)
```

specifies that the fd event f is active in all states where `s=1` and whenever it occurs, the next state is derived from the original one by changing variable `s` to `0` with probability 0.3 and to `2` with probability 0.7. The probabilities in each command have to sum to one. Observe that fd event commands are similar to DTMC commands in PRISM. The synchronization labels are used only to impose impulse rewards as for CTMC, e.g.

```
rewards  [L] true : 1.0;  endrewards
```

The rate rewards are specified the same way as for CTMC in PRISM. The PRISM source code for the fdCTMC of Example 1 is in Appendix B. The implementation details concerning the fdCTMC structure are provided in Appendix B as well.

Users can run the implemented algorithms from both the graphical and the command-line interfaces of PRISM. The expected reward and synthesis implementations are available in menu `Model -> Compute -> Exp. reachability reward` and `Model -> Compute -> FD synthesis`, respectively or using the command-line option `-expreachreward` and `-fdsynthesis`, respectively. The error bound ε is specified in `Options -> Options -> Termination epsilon` or in the command-line option `-epsilon`.

4 Implementation Issues

Implementation of the expected reward computation was a straightforward application of existing PRISM methods. For the synthesis we implemented the *unbounded optimization* algorithm from [1]. The algorithm is based on discretization, i.e. we provide discretization bounds and restrict the uncountable space of delay vectors into a finite space. Instead of an exhaustive search through the finite space, we use the idea of [1] and transform the parametric (discretized) fdCTMC into an MDP where actions correspond to the choices of fd event delays. Now, the minimal solution of the MDP yields the optimal delay vector.

The discretization bounds consist of the discretization step δ, the upper bound on fd event delay $\overline{d}$ and the precision κ for computation of action parameters. They are computed for each fd event separately from the error bound ε, the number of states, the minimal transition probability, and other fdCTMC model attributes. For more detail see Appendix A. Note that in every fdCTMC model, the delays for all fd events have to be specified. Applying these delays, we compute the corresponding expected reward $\overline{Val}$ which is used as an upper bound for the optimal reward. Then $\overline{Val}$ is employed when computing the discretization bounds. The lower the $\overline{Val}$ is, the faster the synthesis implementation performs. Thus it is worth to think of good delays of fd events when specifying the model.

Given the discretization bounds one has to compute the transition probabilities and expected accumulated reward for each action in the MDP corresponding to the discretized delay of fd event. This can be done using the transient analysis of subordinated CTMCs [10].

Prototype Implementation In the first implementation we used straightforward approach to call built-in methods of PRISM to compute the required quantities for each discretized fd event delay separately. This is reasonable since the built-in methods are correctly and efficiently programmed for all PRISM engines and methods of computing transient analysis. However, we experienced that most of the time was spent computing the transient analysis rather than solving the created MDP, e.g. 520 seconds out of 540 seconds of total time.[4] One of the reasons is that in each iteration a small portion of memory is allocated and freed by built-in PRISM methods. Since there is a large number of actions, the amount of reallocated memory was slowing down the computation. Thus we decided to reimplement the computation of transient probabilities the applying principles of dynamic programming.

Iterative Computation of Transient Analysis The transient probabilities can be very efficiently approximated up to an arbitrary small error using the uniformization technique. The problem is that we have to compute the transient probabilities for each value of a very large set $\{i \cdot \delta \mid i \in \mathbb{N}_0 \text{ and } 0 < i \leq \overline{d}/\delta\}$ and allow only fixed error κ for each computation. The transient probability vector $\pi(\delta)$ of a CTMC C at time δ can be computed using uniformization by

$$\pi(\delta) = \sum_{j=0}^{J} \mathbf{1}_{s_{in}} \cdot P^j \cdot \frac{(\lambda \cdot \delta)^j}{j!} \cdot e^{-\lambda \cdot \delta}, \tag{1}$$

where $\mathbf{1}_{s_{in}}$ is the initial vector of C, λ is the uniformization rate of C, and P is the transition kernel of the uniformized C. The choice of number J influences the error of the formula. It is easy to compute the value of J such that the error is sufficiently small.

However, for time $i \cdot \delta$ we can use the previously computed transient probabilities as

$$\pi(i \cdot \delta) = \sum_{j=0}^{J} \pi((i-1) \cdot \delta) \cdot P^j \cdot \frac{(\lambda \cdot \delta)^j}{j!} \cdot e^{-\lambda \cdot \delta}. \tag{2}$$

It is again easy to compute J such that the overall allowed error is not exceeded. Instead of performing naïve computation for each number in $\{i \cdot \delta \mid i \in \mathbb{N}_0 \text{ and } 0 < i \leq \overline{d}/\delta\}$ with according number of steps $J_1, \ldots, J_{\overline{d}/\delta}$ to cause error bounded by κ in each computation, we compute the transient probabilities iteratively with sufficiently large J to cause small error in all computations. For example, if we have $\delta = 0.1$, $\overline{d}/\delta = 1000$, rate $\lambda = 1.0$ and $\kappa = 0.01$ using the naïve method we have to do $J_1 + \cdots + J_{\overline{d}/\delta} = 66,265$ steps and using the iterative method $J \cdot \overline{d}/\delta = 3,000$ steps. This is significant difference since a vector matrix multiplication is performed in each step. Thus we hard-programmed the iterative computation of transient probabilities and accumulated rewards in CTMC what caused a dramatic speedup thanks to the smaller number of arithmetic operations and better memory management.

[4] Computed for the rejuv model and the error bound 0.001, see Section 5.

Precomputation Careful reader may have noticed that (2) can be further simplified to

$$\pi(i \cdot \delta) = \pi((i-1) \cdot \delta) \cdot e^{-\lambda \cdot \delta} \cdot \sum_{j=0}^{J} P^j \cdot \frac{(\lambda \cdot \delta)^j}{j!}. \tag{3}$$

Hence, the matrix $e^{-\lambda \cdot \delta} \cdot \sum_{j=0}^{J} P^j \cdot (\lambda \cdot \delta)^j / j!$ can be easily precomputed beforehand and used for computation of each $\pi(i \cdot \delta)$ to increase the savings even more. However, this is not true. J is small and the matrix P is sparse for the most reasonable models and error bounds. But $e^{-\lambda \cdot \delta} \cdot \sum_{j=0}^{J} P^j \cdot (\lambda \cdot \delta)^j / j!$ is not sparse for almost each error bound, P, and λ, what is known as "fill-in" phenomenon. Thus using (2) is typically more efficient than using (3). Similar observations were discussed in [5].

Implementing the synthesis algorithm of [1], we inherited the following restrictions on the input fdCTMC models. There is at most one concurrently active fd event in each state, i.e. $\forall s \in S \ : \ |A(s)| \leq 1$. For each fd event there is at most one state where its timer is set. Every state has a positive rate reward, i.e. $\forall s \in S \ : \ \mathcal{R}(s) > 0$. Moreover, we add that all fd events have positive impulse rewards, i.e. $\forall f \in F \wedge s, s' \in S \ : N(s,f)(s') > 0 \implies \mathcal{I}(s,f,s') > 0$. For the expected reward implementation only the first two restrictions are valid.

5 Experimental Results

We tested the performance of our synthesis implementation on the model from Example 1 for various sizes of the queue (2, 4, 6, and 8) and the rejuvenation model provided in Appendix B. The considered error bounds are 0.005, 0.0025, 0.0016, 0.00125, and 0.001. The following table shows the expected rewards and the computation times for a given error bound. As the expected rewards are very similar for different error bounds, we show their longest common prefix, instead of listing five similar long numbers.

Model	CPU time [s]					Longest common prefix of exp. rewards
	ε : 0.005	0.0025	0.0016	0.00125	0.00100	
	$1/\varepsilon$: 200	400	600	800	1000	
rejuv	5.87	12.09	14.71	21.60	23.84	0.94431314832
dpm2	58.22	121.15	195.61	234.58	248.52	0.336634754
dpm4	156.02	354.35	509.19	2197.10	2652.05	0.337592724
dpm6	259.76	532.47	2705.45	3026.77	5124.10	0.337583980
dpm8	616.47	3142.44	6362.79	22507.55	27406.62	0.337537611

Note that the computed values of the expected reward are of a much better precision than required. This indicates that there might even be a space for improvements of the synthesis algorithm, e.g. by computation of tighter discretization bounds. It is worth mentioning that the longest computation (dpm8 for error 0.001) took only 1 hour and 30 minutes of real clock time thanks to the native parallelism of Java (the table shows the sum for all threads). Our experiments show that the implementation retains the theoretical complexity bounds saying that the computation time is exponential to the number of states and polynomial to $1/\varepsilon$.

The computations were run on platform HP DL980 G7 with 8 64-bit processors Intel Xeon X7560 2.26GHz (together 64 cores) and 448 GiB DDR3 RAM, but only 304GB was provided to Java. The time was measured by the Linux command `time`.

6 Conclusions and Future Work

In this paper, we incorporated the fdCTMC models into PRISM and implemented the expected reward computation and the synthesis algorithm. The tool is available on `http://www.fi.muni.cz/~xrehak/fdPRISM/`. We have used the explicit state PRISM engine. Based on the promising results, it is reasonable to (re)implement the synthesis and other model checking algorithms for fdCTMCs in the more efficient PRISM engines. Moreover, new effort can be put to reduce the number of current restrictions on the fdCTMC models. For instance the method of stochastic state classes [6] implemented in ORIS may be applied for computation of transient analysis instead of uniformization.

Acknowledgments We thank Vojtěch Forejt and David Parker for fruitful discussions. This work is partly supported by the Czech Science Foundation, grant No. P202/12/G061.

References

1. Brázdil, T., Korenčiak, Ľ., Krčál, J., Novotný, P., Řehák, V.: Optimizing performance of continuous-time stochastic systems using timeout synthesis. In: QEST. LNCS, vol. 9259, pp. 141–159. Springer (2015)
2. Fackrell, M.: Fitting with matrix-exponential distributions. Stochastic models 21(2-3), 377–400 (2005)
3. German, R.: Performance Analysis of Communication Systems with Non-Markovian Stochastic Petri Nets. Wiley (2000)
4. Guet, C., Gupta, A., Henzinger, T., Mateescu, T., Sezgin, A.: Delayed continuous-time Markov chains for genetic regulatory circuits. In: CAV. LNCS, vol. 7358, pp. 294–309. Springer (2012)
5. Haddad, S., Mokdad, L., Moreaux, P.: A new approach to the evaluation of non markovian stochastic Petri nets. In: ICATPN. LNCS, vol. 4024, pp. 221–240. Springer (2006)
6. Horváth, A., Paolieri, M., Ridi, L., Vicario, E.: Transient analysis of non-markovian models using stochastic state classes. Perform. Eval. 69(7-8), 315–335 (2012)
7. Korenčiak, Ľ., Krčál, J., Řehák, V.: Dealing with zero density using piecewise phase-type approximation. In: EPEW. LNCS, vol. 8721, pp. 119–134. Springer (2014)
8. Krčál, J.: Formal Analysis of Discrete-Event Systems with Hard RealTime Bounds. Ph.D. thesis, Faculty of Informatics, Masaryk University, Brno (2014)
9. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: Verification of probabilistic real-time systems. In: CAV. LNCS, vol. 6806, pp. 585–591. Springer (2011)
10. Lindemann, C.: An improved numerical algorithm for calculating steady-state solutions of deterministic and stochastic Petri net models. Perform. Eval. 18(1), 79–95 (1993)
11. Neuts, M.: Matrix-Geometric Solutions in Stochastic Models: An Algorithmic Approach. The Johns Hopkins University Press, Baltimore (1981)
12. Puterman, M.: Markov Decision Processes. Wiley (1994)
13. Qiu, Q., Wu, Q., Pedram, M.: Stochastic modeling of a power-managed system: construction and optimization. In: ISLPED. pp. 194–199. ACM Press (1999)
14. Xie, W., Sun, H., Cao, Y., Trivedi, K.S.: Optimal webserver session timeout settings for web users. In: Computer Measurement Group Conferenceries. pp. 799–820 (2002)
15. Zimmermann, A.: Applied restart estimation of general reward measures. In: RESIM. pp. 196–204 (2006)
16. Zimmermann, A.: Modeling and evaluation of stochastic Petri nets with timenet 4.1. In: ICST. pp. 54–63. IEEE (2012)

A Discretization Bounds

Using the full version of [1] we derived the exact formulas of the discretization bounds for each fd event:

$$\overline{d} = \max\Big\{\frac{\overline{Val}}{minP^{|S_{fd}|} \cdot minR}\ ;\ \frac{e \cdot |\ln(\alpha/2)|}{\lambda \cdot minP}\Big\},$$
$$\delta = \frac{\alpha}{D_1},$$
$$\kappa = \frac{\varepsilon \cdot \delta \cdot minR}{2 \cdot |S'| \cdot (1 + \overline{Val})},$$

where

$$\alpha = \min\Big\{\frac{\varepsilon}{Bound[\#] \cdot (1 + \overline{Val}) \cdot |S'|}\ ;\ \frac{1}{2 \cdot Bound[\#] \cdot |S'|}\Big\},$$
$$D_1 = \max\{2 \cdot \lambda\ ;\ 1 \cdot (\lambda + 1) \cdot maxR\},$$

- $Bound[\#]$ is an upper bound on expected number of steps to reach target from any state in the created MDP, i.e.

$$Bound[\#] = \frac{\overline{Val}}{\text{minimal expected one-step reward in the created MDP}},$$

- $\overline{Val}$ is the upper bound on the expected reward,
- S' is the state space of the created MDP,
- λ is the uniformization rate, and
- $|S_{fd}|$, $minP$, $maxR$, and $minR$ is the number of states, the minimal branching probability, the maximal reward, and the minimal reward in the subordinated CTMC for the given fd event, respectively.

B fdCTMC in PRISM Language

Extension of PRISM data structures The `FDCTMCSimple` class extends the `CTMCSimple` class by a vector of objects of type `FDEvent` and few methods to work with them (the methods are explained in the corresponding interface `FDCTMC` that is an extension of interface `CTMC`). The `FDEvent` is basically an extension of `DTMCSimple` class by one `double` attribute that keeps delay of the fd event and one `String` attribute that keeps the label of the fd event. The transition kernel is kept in the inherited attributes from the `DTMCSimple` class.

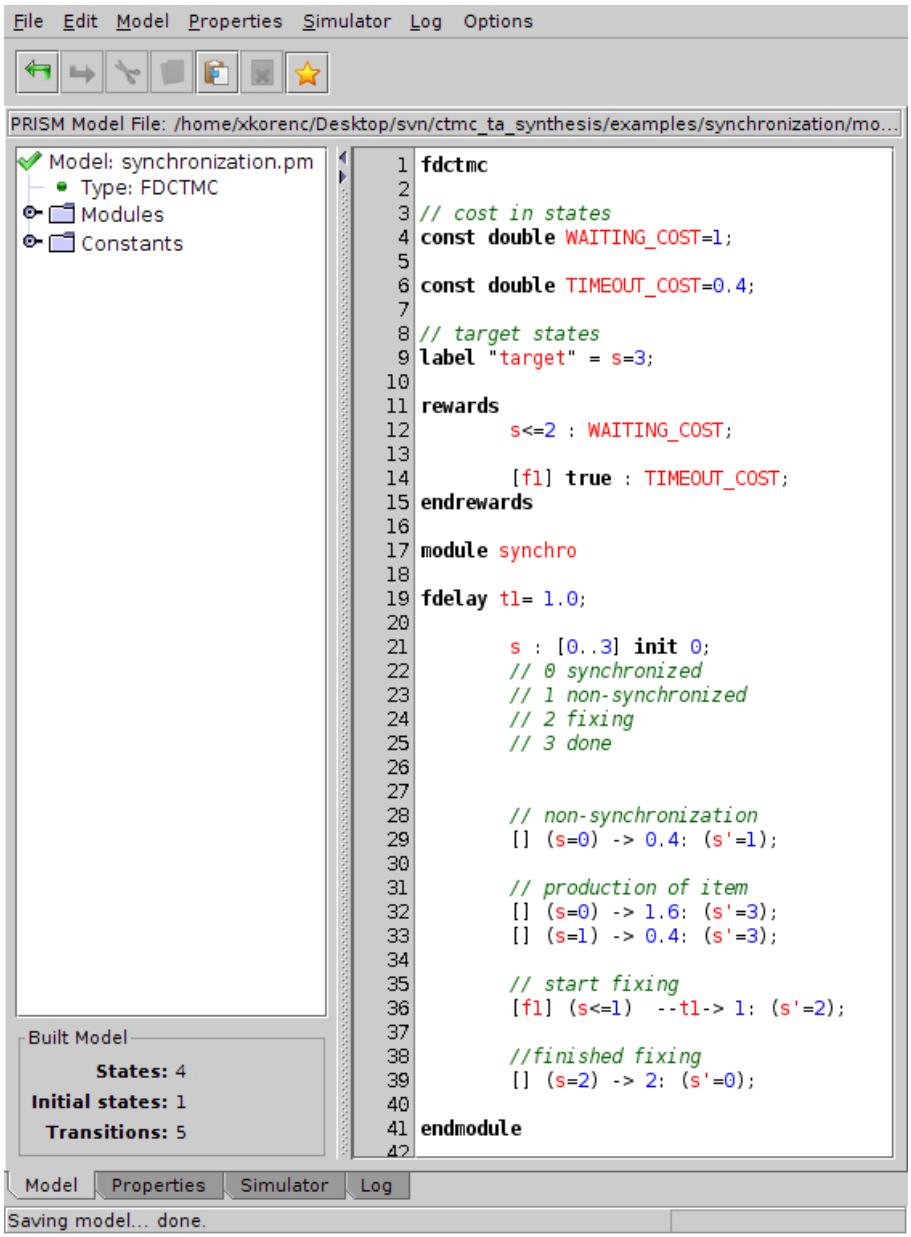

Fig. 1. The graphical user interface of PRISM with the source code of the rejuvenation model [3].

```
fdctmc

// Maximum queue size
const int q_max = 2;

// target states
label "target" = sp=3;

// Request arrival rate (mean inter-arrival time is 0.72 seconds)
const double rate_arrive = 1/0.72;

// Rate of service (average service time = 0.08s)
const double rate_serve = 1/0.08;

rewards
        sp=3 : 0.35;
        sp=2 : 0.95;
        sp=1 : 0.35;
        sp=0 : 0.13;

        [sleep2idle] true : 0.067;
        [idle2sleep] true : 0.006;
endrewards

module disk_drive

// time to sleep mode
fdelay f1= 1.0;

// time of sleeping
fdelay f2= 2.0;

        // q = number of requests currently in queue
        q : [0..q_max] init 0;

        // Power state of SP: 0=sleep, 1=idle, 2=busy, 3=target
        sp : [0..3] init 1;

        // arrival of task
        [e1] (sp=1)          -> rate_arrive : (sp'=2) & (q'=min(q+1,q_max));
        [e1] (sp=0) | (sp=2) -> rate_arrive : (q'=min(q+1,q_max));

        // A request is served
        [serve] (sp=2) & (q>1) -> rate_serve:(q'=q-1);
        // Last request is served
        [serve] (sp=2) & (q=1) -> rate_serve:(sp'=3);

        // Respond to controls from power manager (PM):
        // Switch from idle mode to sleep mode
        [idle2sleep] sp=1 & q=0 --f1-> (sp'=0);

        // Switch from sleep mode to idle mode
        [sleep2idle] sp=0 & q=0 --f2-> (sp'=0);
        [sleep2idle] sp=0 & q>0 --f2-> (sp'=2);

endmodule
```

Fig. 2. The source code of the fdCTMC from Example 1 in the PRISM language.